

A qubit-efficient solver for the Low Autocorrelation Binary Sequences problem

Marco Sciorilli,^{1,2} Giancarlo Camilo,¹ Thiago O. Maciel,¹ Askery Canabarro,^{1,3} Lucas Borges,^{1,4} and Leandro Aolita¹

¹*Quantum Research Center, Technology Innovation Institute, Abu Dhabi, UAE*

²*Laboratoire d'Informatique de Paris 6, CNRS,
Sorbonne Université, 4 Place Jussieu, 75005 Paris, France*

³*Federal University of Alagoas, Campus Arapiraca, Arapiraca, AL 57309-005, Brazil*

⁴*Federal University of Rio de Janeiro, Caixa Postal 652, Rio de Janeiro, RJ 21941-972, Brazil*

Pauli Correlation Encoding (PCE) is a qubit-efficient variational approach to combinatorial optimizations. It offers a polynomial reduction in qubit count and a super-polynomial suppression of barren plateaus. We extend the PCE-based framework to the Low Autocorrelation Binary Sequences (LABS) problem, a notoriously hard benchmarking problem for classical and quantum solvers. We simulate the solver for LABS instances of up to $N = 45$ variables using only 4 qubits and a circuit Ansatz with 30 two-qubit gates. We find that encodings based on anti-commuting Paulis strings attain higher solution-qualities than those with commuting strings, hinting at a possible connection between quantum correlations and solver performance. Moreover, our method displays a better scaling in time to solution than quantum schemes such as QAOA and even the best classical heuristic, Tabu search, while using only a fraction of the quantum resources required by other quantum approaches. Finally, we perform loss-function evaluations on IonQ's **Forte** experimental quantum processor for $N = 120$ using only 10 qubits, showcasing good noise reliance even without error mitigation. Our findings highlight the potential of PCE as a promising tool to reduce the resource requirements for quantum optimization heuristics.

Main idea. The present work investigates the use of Pauli correlation encoding (PCE)[1] to solve LABS on NISQ devices, focusing in particular on how the algebraic structure of the Pauli set affects performance. In PCE, the LABS variables are encoded as signs of Pauli expectation values, $x_i = \text{sgn}(\langle \Pi_i \rangle)$, so that the choice of Π becomes a central component. We compare set of commuting paulis, $\Pi^{(C)}$, which define a jointly measurable and more classical encoding, with set of non-commuting ones, $\Pi^{(NC)}$, which probe complementary features of the quantum state. Numerical simulations show that solver performance improves with the fraction of non-commuting relations, and that $\Pi^{(NC)}$ yields better time-to-solution behavior. We therefore adopt this non-commuting encoding as the main PCE construction for benchmarking against QAOA and Tabu Search, and observe in numerical simulations that the PCE-based LABS solver exhibits improved exponential runtime scaling compared to QAOA while remaining highly competitive with state-of-the-art classical heuristics such as Tabu search [2].

The LABS problem. Given an integer $N > 2$, the Low Autocorrelation Binary Sequences (LABS) problem consists of finding a binary sequence $x \in \{-1, +1\}^N$ that minimizes the sidelobe energy

$$E(x) = \sum_{k=1}^{N-1} C_k^2(x), \quad C_k(x) = \sum_{i=1}^{N-k} x_i x_{i+k}, \quad k \in [N-1], \quad (1)$$

where $C_k(x)$ is the k -th autocorrelation of the sequence. Equivalently, one may maximize the merit factor $F(x) = N^2/[2E(x)]$, which measures the peak-to-sidelobe ratio and is relevant in applications such as communications, radar, and cryptography [3–5].

LABS is NP-hard and provides a hard benchmark for optimization algorithms. Its landscape has a characteristic “golf-course” structure, with many local minima and highly isolated global optima [6]. Exact enumeration requires searching over 2^N sequences, while the best classical heuristics, including Tabu-based methods, exhibit empirical exponential time-to-solution scaling [7, 8]. Also, for $N > 200$, the performance of classical heuristics degrades significantly. Recent quantum approaches based on QAOA and quantum minimum finding suggest improved asymptotic scaling, but require resources beyond near-term devices [9]. This makes LABS a natural testbed for qubit-efficient variational heuristics such as Pauli Correlation Encoding.

A qubit-efficient quantum LABS solver. Here we propose a qubit-efficient solver for the LABS problem following the *Pauli-correlation encoding* (PCE) framework recently introduced in [1]. Given a set $\Pi \equiv \{\Pi_i\}_{i \in [N]}$ of n -qubit Pauli strings Π_i , PCE encodes binary variables x_i into correlators of this set as

$$x_i := \text{sgn}(\langle \Pi_i \rangle), \quad i \in [N], \quad (2)$$

where sgn is the sign function and $\langle \Pi_i \rangle := \langle \Psi | \Pi_i | \Psi \rangle$ denotes the expectation value of Π_i on a quantum state $|\Psi\rangle$. We take $|\Psi\rangle$ as the output of a parameterized quantum circuit with a brickwork architecture and parameters θ to be classically




FIG. 1: **Left.** Average approximation ratio over 1000 random initializations for LABS sizes $N \in \{27, 43, 87, 173\}$ as a function of the ratio of non-commuting to commuting relations among the Pauli operators in the encoding set. The trend shows improved average solution quality as non-commutativity increases. **Right.** TTS scaling comparison for the PCE LABS solver using Pauli sets sampled from $\Pi^{(NC)}$ and $\Pi^{(C)}$ for $N \in [28, 36]$. TTS is the number of cost-function evaluations required to reach the exact solution [8]. The non-commuting construction yields a better empirical exponential basis, approximately $\mathcal{O}(1.345^N)$ versus $\mathcal{O}(1.407^N)$ for the commuting construction.

optimized using a variational approach [1]. The goal of the parameter optimization is to minimize the non-linear loss function

$$\mathcal{L} = \sum_{\ell=1}^{N-1} \sum_{i,j=1}^{N-\ell} \tilde{x}_i \tilde{x}_{i+\ell} \tilde{x}_j \tilde{x}_{j+\ell} - \beta \sum_{i=1}^N \tilde{x}_i^2, \quad (3)$$

where $\tilde{x}_i := \tanh(\alpha \langle \Pi_i \rangle)$ is a real-valued relaxation of the binary variable to make it better-suited for gradient-based optimization. The rescaling factor $\alpha > 1$ is introduced to restore the non-linear behaviour of the tanh, given that the n qubit correlators have magnitude decreasing polynomially with n . The last term in (3) is a regularization term forcing all correlators away from zero, a strong trivial attractor for the relaxed energy landscape, which was observed to improve the solver's performance. Once the training is complete, the circuit output state is measured and a bit-string x is obtained via Eq. (2).

We studied the performance of PCE with two different choices for the Pauli correlator set Π . The first, $\Pi^{(C)}$, is the same set used in [1] of k -body correlators of the form $X^k \mathbb{1}^{n-k}, Y^k \mathbb{1}^{n-k}, Z^k \mathbb{1}^{n-k}$ (and permutations thereof). This is the union of 3 sets of mutually-commuting Pauli strings and corresponds to a polynomial compression of degree k since there are $3 \binom{n}{k}$ such correlators. The second, $\Pi^{(NC)}$, is a set of maximally non-commuting Pauli strings constructed when finding the maximal commuting basis of orthogonal unitary matrices using the Pauli group [10]. The intuition is that using measurements that are maximally non-commutative, the outcomes we get from one of these Pauli operators leaves us completely uncertain about the outcomes of measurements over another [11]. This increase the independence and therefore the expressivity of the assigned variables x_i in Eq. (2). We ran numerical simulations in which the number of qubits, variational ansatz, and optimizer stopping criterion were fixed while both the ansatz initialization and the Pauli set were randomized. The average approximation ratio increases with the fraction of non-commuting pairwise relations inside the encoding set, showing a clear positive correlation between non-commutativity and solution quality. Moreover, a direct comparison of the two constructions shows that $\Pi^{(NC)}$ leads to better TTS scaling than $\Pi^{(C)}$, which is why it is used as the main figure of merit in the benchmark against other solvers. The results are shown in Fig. 1.

Numerical results. We perform numerical simulations using the two PCE-based quantum solvers for LABS instances of sizes $N \leq 45$ using the Pauli-correlation set $\Pi^{(NC)}$ with $n = 4$ qubits. To determine a suitable circuit depth, we increased

Heuristic		TTS	CI	R^2	Heuristic		TTS	CI	R^2
QAOA [9]		1.46	[1.42, 1.50]	>0.94	QAOA [9]		1.46	[1.42, 1.50]	>0.94
Tabu Search	Median	1.358	[1.297, 1.422]	0.97	Tabu Search	Median	1.413	[1.334, 1.496]	0.99
	Ensemble	1.370	[1.355, 1.386]	0.78		Ensemble	1.404	[1.392, 1.417]	0.90
PCE	Median	1.331	[1.276, 1.387]	0.95	PCE	Median	1.309	[1.270, 1.349]	0.99
	Ensemble	1.335	[1.315, 1.355]	0.75		Ensemble	1.313	[1.297, 1.330]	0.93
PCE + Mem. Tabu	Median	1.245	[1.232, 1.255]	0.98	PCE + Mem. Tabu	Median	1.319	[1.309, 1.329]	0.99

TABLE I: Estimated base value b , controlling the exponential scaling $\mathcal{O}(b^N)$, and associated R^2 values for the TTS linear fit across four different algorithms - PCE, PCE + Memetic Tabu Search, QAOA, and Tabu Search - for even (left) and odd (right) instances.

the number of parameterized layers for small problem instances until no significant improvement in the average approximation ratio over 1000 runs was observed.

We observe that the optimal depth scales linearly with n , and therefore sublinearly with N . In the reported benchmarks, we fix the number of two-qubit gates to 30, corresponding to 15 layers of the brickwork circuit and a total of 150 parameters.

As our figure of merit, we adopt the time-to-solution (TTS), defined as the number of cost-function evaluations required to observe the exact solution, following [8]. Although the time required to reach optimality is a rigorous metric, To obtain empirical scalings, we use problem instances with unique canonical solutions, *i.e.*, cases where the ground state is only degenerate due to the inherent LABS symmetries [8]. For each instance size, we run the optimization until getting the optimum 50 times, except for $N \in [41, 45]$, where we run until 10 found optimum. We fit the exponential curve $TTS = c \cdot b^N$ by linear regression on $\log(TTS)$, using both the median TTS for each sequence length and the full ensemble of data points.

The results show a clear scaling advantage of PCE compared with QAOA for both even and odd instances, reported in Table I. In addition, we compare with Tabu Search without its memetic component, which was observed not to provide a scaling advantage [8], and we choose the Tabu tenure parameters from the recent GPU implementation in [12]. The PCE results are competitive within error bars. While minor polynomial advantage factors do not alter the overall exponential scaling, even small improvements in the exponential basis can become important for larger problem instances. Running PCE on quantum hardware also requires accounting for the overhead of gate operations, repeated circuit executions to estimate the loss function and gradients to precision ϵ , and gradient estimation. We estimate that the gate cost scales as $\mathcal{O}(N)$, the sample complexity as $\mathcal{O}(N^7/\epsilon^2)$, and the gradient-estimation overhead as $\mathcal{O}(N)$, giving an overall hardware overhead of $\mathcal{O}(N^9/\epsilon^2)$. Even with this polynomial overhead, extrapolating the observed TTS scalings suggests a possible quantum-advantage crossover at instance sizes of order thousands, corresponding in the PCE mapping to fewer than 10^2 qubits and two-qubit gate counts of order 10^3 . Approximate advantages may appear earlier, as classical exact solutions are known only up to limited sizes.

Experimental deployment. We experimentally demonstrate our quantum solver on IonQ’s **Forte** trapped-ion quantum processor using a LABS instance of size $N = 120$. For this proof-of-principle deployment, we use the quadratic Pauli-correlation encoding and restrict to the same-letter Pauli observables XX , YY , and ZZ used in [1]; this hardware-friendly choice reduces the measurement overhead to three bases. The circuit is a standard VQA brickwork ansatz expressed in the hardware-native gate set, consisting of alternating layers of parametrized entangling RZZ gates and single-qubit rotation layers built from the native GPI and GPI2 operations. Its depth is fixed to 10 in the hardware-oriented deployment setup. Circuit parameters are optimized offline by classical simulation, and the resulting pre-trained circuit is then executed on the hardware. For three independent random initializations of the ansatz parameters, we measure the loss function at both the first and last optimization epochs, accumulating measurement statistics until the estimated loss values no longer changed appreciably. The experimental data, reported in Fig. 2, show that shot noise has only a minor impact on the solution quality: a few thousand shots already provide sufficient statistics for estimating the loss during training. Hardware noise, in contrast, does not prevent training but noticeably degrades the final solution quality and can erode the expected TTS scaling advantage. At the tested circuit depth, barren-plateau effects should be modest; nevertheless, the relative error is larger at initialization because the Pauli expectation values are close to zero, making the measured signal more sensitive to hardware and finite-shot fluctuations. As optimization progresses, the expectation values move away from zero, the signal-to-noise ratio improves, and the relative error decreases. Finally, at the last optimization epoch, the rounding step that maps the relaxed experimental solution back to binary values yields a bitstring error rate of only 1.6%, indicating that hardware errors do not meaningfully alter the final binary solution in this deployment.

Discussion. The solver we propose achieves competitive TTS scaling using only $n = 4$ qubits and 30 two-qubit gates. Our results suggest that this performance depends not only on the compression rate, but also on the choice of Pauli correlation set II. Although PCE-based variational methods have no performance guarantees, their performance on MaxCut and LABS motivates further study. The trapped-ion experiment on an $N = 120$ LABS instance suggests that useful loss estimates can be obtained with moderate shot budgets, and that the final rounded binary solutions can tolerate some hardware noise. For the quadratic-compression variant used here in the hardware deployment, $n = 8$ qubits would cover the largest LABS instances with known exact solutions, $N = 66$, while $n \leq 20$ qubits would reach the largest instances currently tackled by classical heuristics [13]. Altogether, these regimes open several exciting directions for future work: benchmarking PCE as a quantum-inspired heuristic, exploring its use as a warm start for classical solvers [14], and assessing more broadly whether compressed variational methods can offer practical advantages.




FIG. 2: **Experiment.** *Main panels.* Relative mean-squared error (RMSE) of the measured loss values for three different parameter initializations, evaluated at the first optimization epoch (left) and at epoch 600, the final one (right), as a function of the number of measurement shots. Experimental data were collected on the IonQ trapped-ion Forte-1 QPU. For comparison, we also report the corresponding shot-noise-only RMSE obtained from numerical simulations. The circuit uses 10 qubits and 30 two-qubit gates, targeting a LABS instance of size $N = 120$. *Insets.* Box-plot distributions of the mean approximation ratio obtained over 500 random initializations by simulating the full optimization under three noise models: exact state-vector simulation (Exact), shot-noise-only (Shot noise), and hardware-noise matching the observed QPU error rates (QPU noise).

-
- [1] M. Sciorilli, L. Borges, T. L. Patti, D. García-Martín, G. Camilo, A. Anandkumar, and L. Aolita, Towards large-scale quantum optimization solvers with few qubits, *Nature Communications* **16**, 476 (2025).
 - [2] F. Glover, Tabu search—part i, *ORSA Journal on Computing* **1**, 190 (1989).
 - [3] S. W. Golomb and R. A. Scholtz, Generalized barker sequences, *IEEE Transactions on Information Theory* **11**, 533 (1965).
 - [4] T. Packebusch and S. Mertens, Low autocorrelation binary sequences, *Journal of Physics A: Mathematical and Theoretical* **49**, 165001 (2016).
 - [5] I. P. Gent and T. Walsh, CSPLib: A Benchmark Library for Constraints (1999), available online at <http://www.csplib.org/> as prob005.
 - [6] S. Mertens and C. Bessenrodt, On the ground states of the bernasconi model, *Journal of Physics A: Mathematical and General* **31**, 3731–3749 (1998).
 - [7] J. E. Gallardo, C. Cotta, and A. J. Fernández, Finding low autocorrelation binary sequences with memetic algorithms, *Applied Soft Computing* **9**, 1252 (2009).
 - [8] B. Bošković, F. Brglez, and J. Brest, Low-autocorrelation binary sequences: On improved merit factors and runtime predictions to achieve them, *Applied Soft Computing* **56**, 262 (2017).
 - [9] R. Shaydulin, C. Li, S. Chakrabarti, M. DeCross, D. Herman, N. Kumar, J. Larson, D. Lykov, P. Minssen, Y. Sun, Y. Alexeev, J. M. Dreiling, J. P. Gaebler, T. M. Gatterman, J. A. Gerber, K. Gilmore, D. Gresh, N. Hewitt, C. V. Horst, S. Hu, J. Johansen, M. Matheny, T. Mengle, M. Mills, S. A. Moses, B. Neyenhuis, P. Siegfried, R. Yalovetzky, and M. Pistoia, Evidence of scaling advantage for the quantum approximate optimization algorithm on a classically intractable problem, *Science Advances* **10**, 10.1126/sciadv.adm6761 (2024).
 - [10] Bandyopadhyay, Boykin, Roychowdhury, and Vatan, A new proof for the existence of mutually unbiased bases, *Algorithmica* **34**, 512 (2002).
 - [11] W. K. Wootters and B. D. Fields, Optimal state-determination by mutually unbiased measurements, *Annals of Physics* **191**, 363 (1989).
 - [12] Z. Zhang, J. Shen, N. Kumar, and M. Pistoia, New improvements in solving large labs instances using massively parallelizable memetic tabu search (2025), arXiv:2504.00987 [cs.DC].
 - [13] B. Pšeničnik, R. Mlinarič, J. Brest, and B. Bošković, Dual-step optimization for binary sequences with high merit factors, *Digital Signal Processing* **165**, 105316 (2025).
 - [14] A. G. Cadavid, P. Chandarana, S. V. Romero, J. Trautmann, E. Solano, T. L. Patti, and N. N. Hegade, Scaling advantage with quantum-enhanced memetic tabu search for labs (2025), arXiv:2511.04553 [quant-ph].